arXiv:1607.05597v1 [cs.DS] 19 Jul 2016

Distributed Construction of Purely Additive Spanners*

Keren Censor-Hillel† Telikepalli Kavitha‡ Ami Paz† Amir Yehudayoff§

July 20, 2016

Abstract

This paper studies the complexity of distributed construction of purely additive spanners in the CONGEST model. We describe algorithms for building such spanners in several cases. Because of the need to simultaneously make decisions at far apart locations, the algorithms use additional mechanisms compared to their sequential counterparts.

We complement our algorithms with a lower bound on the number of rounds required for computing pairwise spanners. The standard reductions from set-disjointness and equality seem unsuitable for this task because no specific edge needs to be removed from the graph. Instead, to obtain our lower bound, we define a new communication complexity problem that reduces to computing a sparse spanner, and prove a lower bound on its communication complexity using information theory. This technique significantly extends the current toolbox used for obtaining lower bounds for the CONGEST model, and we believe it may find additional applications.

1 Introduction

A graph *spanner* is a sparse subgraph that guarantees some bound on how much the original distances are stretched. Graph spanners, introduced in 1989 [PS89, PU89a], are fundamental graph structures which are central for many applications, such as synchronizing distributed networks [PU89a], information dissemination [CHKM12], compact routing schemes [Che13a, PU89b, TZ01], and more.

Due to the importance of spanners, the trade-offs between their possible sparsity and stretch have been the focus of a huge amount of literature. Moreover, finding time-efficient constructions of spanners with optimal guarantees has been a major goal for the distributed computing community, with ingenious algorithms given in many studies (see, e.g., [EP04, Elk05, EZ06, DGPV09, BKMP10, BS07, DG08, DGP07, DGPV08, DMP+05, Pet10]). One particular type of spanners are *purely additive spanners*, in which the distances are promised to be stretched by no more than an additive term. However, distributed constructions of such spanners have been scarce, with the only known construction being a $(+2)$-additive spanner construction with $O(n^{3/2}\log n)$ edges in $O(\sqrt{n}\log n + D)$ rounds in a network of size n and diameter D [LP13] (also follows from [HW12]).

The absence of distributed constructions of purely additive spanners is explicitly brought into light by Pettie [Pet10], and implicitly mentioned in [DGP07].

†Department of Computer Science, Technion, Israel. {ckeren,amipaz}@cs.technion.ac.il. Supported by ISF individual research grant 1696/14. Part of this work was done while Ami Paz was visiting TIFR, Mumbai.

‡Tata Institute of Fundamental Research, India. kavitha@tcs.tifr.res.in.

§Department of Mathematics, Technion, Israel. amir.yehudayoff@gmail.com.

*An extended abstract of this work will be presented in DISC 2016.

This paper remedies this state of affairs, by providing a study of the complexity of constructing sparse purely additive spanners in the synchronous CONGEST model [Pel00], in which each of n nodes can send an $O(\log n)$-bit message to each of its neighbors in every round. Our contribution is twofold: first, we provide efficient constructions of several spanners with different guarantees, and second, we present new lower bounds for the number of rounds required for such constructions, using tools that are not standard in this context.

1.1 The Challenge

A subgraph H of an undirected unweighted graph $G = (V, E)$ is called a purely additive spanner with stretch β if for every every pair $(u, v) \in V \times V$, we have $\delta_H(u, v) \leq \delta_G(u, v) + \beta$, where $\delta_H(u, v)$ is the u-v distance in H and $\delta_G(u, v)$ is the u-v distance in G. The goal in spanner problems is to construct a subgraph H that is as sparse as possible with β as small as possible, i.e., we seek a sparse subgraph of G which approximates all distances with a small stretch.

The problem of computing sparse spanners with small stretch β is well-studied and we know how to construct sparse purely additive spanners for $\beta = 2, 4, 6$. These have sizes $O(n^{3/2})$ [ACIM99], $\tilde{O}(n^{7/5})$ [Che13b], and $O(n^{4/3})$ [BKMP10], respectively. In a very recent breakthrough, it was shown that there is *no* purely additive spanner of size at most $n^{4/3}/2^{O(\sqrt{\log n})}$ [AB16a].

In a bid to get sparser subgraphs than all-pairs spanners with the same stretch, the following relaxation of *pairwise spanners* has attracted recent interest. Here we are given $\mathcal{P} \subseteq V \times V$: these are our "relevant pairs" and we seek a sparse subgraph which approximates distances between all pairs in $\mathcal{P}$ with a small stretch. That is, for every pair $(u, v) \in \mathcal{P}$, the graph H should satisfy $\delta_H(u, v) \leq \delta_G(u, v) + \beta$ and for pairs (u, v) outside $\mathcal{P}$, the value $\delta_H(u, v)$ could be arbitrarily large. Such a subgraph H is called a $(+\beta)$-pairwise spanner. We use $\tau(\mathcal{P})$ to denote the number of nodes appearing in $\mathcal{P}$, i.e. $\tau(\mathcal{P}) = |\{u \mid \exists v : \{u, v\} \in \mathcal{P}\}|$.

The problem of constructing sparse pairwise spanners was first studied by Coppersmith and Elkin [CE06] who showed sparse subgraphs where distances for pairs in $\mathcal{P}$ were *exactly preserved*; these subgraphs were called *pairwise preservers*. A natural case for $\mathcal{P}$ is when $\mathcal{P} = S \times V$, where $S \subseteq V$ is a set of *source* nodes — here we seek for a sparse subgraph that well-approximates s-v distances for all $(s, v) \in S \times V$. Such pairwise spanners are called *sourcewise spanners*. Another natural setting is when $\mathcal{P} = S \times S$ and such pairwise spanners are called *subsetwise spanners*.

Purely additive spanners are usually built in three steps: first, building clusters which contain all high-degree nodes and adding all the edges of the unclustered nodes; second, building BFS trees which (+2)-approximate all the paths with many missing edges; and third, adding more edges to approximate the other paths.

While our constructions follow the general outline of known sequential constructions of pairwise additive spanners [Kav15, KV13], their techniques cannot be directly implemented in a distributed setting. In the sequential setting, the *clustering* phase is implemented by repeatedly choosing a high-degree node and adding some of its edges to the spanner; these neighbors are marked and ignored in the rest of the phase. In the distributed setting, going over high degree nodes one by one, creating clusters and updating the degrees is too costly. Instead, we choose the cluster centers at random, as done by Thorup and Zwick [TZ05], Baswana and Sen [BS07], and Chechik [Che13b] (see also Aingworth et al. [ACIM99] for an earlier use of randomization for the a dominating set problem).

Sources for BFS trees are carefully chosen in the sequential setting by approximately solving a set-cover problem, in order to cover all paths with many missing edges. Once again, this cannot

be directly implemented in the distributed setting, as the knowledge of all paths cannot be quickly gathered in one location, so we choose the BFS sources randomly [Che13b]. In both the clustering and BFS phases, the number of edges increases by a multiplicative $\log^c n$ factor, for $c < 1$.

The main challenge left is to choose additional edges to add to the spanner in order to approximate the remaining paths well. To this end, we make heavy use of the parallel-BFS technique of Holzer and Wattenhofer [HW12], which allows to construct BFS trees rooted at s different nodes in $O(s + D)$ rounds. We use this technique to count edges in a path, to count missing edges in it, and to choose which edges to add to the spanner. Yet, interestingly, we are unable to match the guarantee on the number of edges of more sophisticated algorithms [BKMP10, Kav15, Woo10]. Some of these algorithms use the *value* of a path, which is roughly the number of pairs of cluster that get closer if the path is added to the spanner. We are not able to measure this quantity efficiently in the distributed setting, and this is one of the reasons we are unable to introduce (+6)-all-pairs spanner matching the sequential constructions.

1.2 Our Contribution

We provide various spanner constructions in the CONGEST model, as summarized in Tables 1 and 2.

Spanner Type	Number of Edges — Distributed		Number of Edges — Sequential	
(+2)-sourcewise	$O\left(n^{5/4}\left\|S\right\|^{1/4}\log^{3/4} n\right)$	(Thm.1)	$O\left(n^{5/4}\left\|S\right\|^{1/4}\log^{1/4} n\right)$	[KV13]
(+2)-pairwise	$O\left(n\left\|\mathcal{P}\right\|^{1/3}\log^{2/3} n\right)$	(Thm.9)	$O\left(n\left\|\mathcal{P}\right\|^{1/3}\right)$	[AB16b]
(+4)-pairwise	$O\left(n\left\|\mathcal{P}\right\|^{2/7}\log^{6/7} n\right)$	(Thm.13)	$O\left(n\left\|\mathcal{P}\right\|^{2/7}\log^{3/7} n\right)$	[Kav15]
(+4)-all-pairs	$O\left(n^{7/5}\log^{4/5} n\right)$	(Thm.5)	$O\left(n^{7/5}\log^{1/5} n\right)$	[Che13b]
(+8)-all-pairs	$O\left(n^{15/11}\log^{10/11} n\right)$	(Thm.18)	$O\left(n^{4/3}\right)$	[BKMP10]
(+2)-subsetwise	$O\left(n\left\|S\right\|^{2/3}\log^{2/3} n\right)$	(Cor.17)	$O\left(n\left\|S\right\|^{1/2}\right)$	[Pet09, CGK13]
(+4)-subsetwise	$O\left(n\left\|S\right\|^{4/7}\log^{6/7} n\right)$	(Cor.17)	$O\left(n\left\|S\right\|^{1/2}\right)$	[Pet09, CGK13]

Table 1: The number of edges in our new constructions versus prior, sequential work. We compare our (+4)-subsetwise with a sequential construction of a (+2)-subsetwise spanner, and our (+8)-all-pairs spanner with a (+6)-all-pairs spanner.

The distributed spanner construction algorithms we present have three main properties: stretch, number of edges, and running time. All three properties hold w.h.p. (with high probability). That is, the algorithm stops in the desired time, with the desired number of edges and the spanner produced has the desired stretch with probability $1 - n^{-c}$, where c is constant of choice. However, we can trade the properties and guarantee two of the three to always hold: if the spanner is too dense or the stretch is too large, we can repeat the algorithm; if the running time exceeds some threshold, we can stop the execution and output the whole graph to get 0 stretch, or output an empty graph to get the desired number of edges. The edges of the constructed spanner can be counted over a BFS tree in G within $O(D)$ rounds. In sourcewise, subsetwise and pairwise spanners, the stretch is measured by running BFS from the relevant nodes (nodes of S of appearing in $\mathcal{P}$) for $O(D)$ rounds in G and again in H; in all-pairs spanners, the stretch is measured by measuring the stretch of the underlying sourcewise or subsetwise spanner.

Spanner Type	Number of Rounds		Lower Bounds	
(+2)-sourcewise	$O\left(\lvert S\rvert + D\right)$	(Thm.1)	$\min\left\{\tilde{\Omega}\left(\frac{n^{3/8}}{\lvert S\rvert^{1/8}}\right), \Omega\left(D\right)\right\}$	[Pet10]
(+2)-pairwise	$O\left(\tau(\mathcal{P}) + D\right)$	(Thm.9)	$\min\left\{\tilde{\Omega}\left(\frac{n^{1/2}}{\lvert\mathcal{P}\rvert^{1/6}}\right), \Omega\left(D\right)\right\}$	[Pet10]
			$\Omega\left(\frac{\lvert\mathcal{P}\rvert}{n\log n}\right)$	(Thm.23)
(+4)-pairwise	$O\left(\tau(\mathcal{P}) + D\right)$	(Thm.13)	$\min\left\{\tilde{\Omega}\left(\frac{n^{1/2}}{\lvert\mathcal{P}\rvert^{1/7}}\right), \Omega\left(D\right)\right\}$	[Pet10]
			$\Omega\left(\frac{\lvert\mathcal{P}\rvert}{n\log n}\right)$	(Thm.26)
(+4)-all-pairs	$O(n^{3/5}\log^{1/5} n + D)$	(Thm.5)	$\min\left\{\tilde{\Omega}\left(n^{3/10}\right), \Omega\left(D\right)\right\}$	[Pet10]
(+8)-all-pairs	$O(n^{7/11}\log^{1/11} n + D)$	(Thm.18)	$\min\left\{\tilde{\Omega}\left(n^{7/22}\right), \Omega\left(D\right)\right\}$	[Pet10]
(+2)-subsetwise	$O(\lvert S\rvert + D)$	(Cor.17)	$\min\left\{\tilde{\Omega}\left(\frac{n^{1/2}}{\lvert S\rvert^{1/3}}\right), \Omega\left(D\right)\right\}$	[Pet10]
(+4)-subsetwise	$O(\lvert S\rvert + D)$	(Cor.17)	$\min\left\{\tilde{\Omega}\left(\frac{n^{1/2}}{\lvert S\rvert^{2/7}}\right), \Omega\left(D\right)\right\}$	[Pet10]

Table 2: Running time: algorithms versus lower bounds, for number of edges as in Table 1. $\tilde{\Omega}$ hides polylogarithmic factors.

We complement our algorithms with some lower bounds for the CONGEST model. We show that any algorithm that constructs an additive (+2)-pairwise spanner with m edges on $p \leq m$ pairs must have at least $\Omega(p/(n\log n))$ rounds, as long as $m \leq n^{3/2}$. For example, a CONGEST construction of a (+2)-pairwise spanner must take $\tilde{\Omega}(\sqrt{n})$ rounds. We also prove lower bounds for (α,β)-pairwise spanners (i.e., for which $\delta_H(u,v) \leq \alpha\,\delta_G(u,v)+\beta$). We show that any algorithm that constructs an (α,β)-pairwise spanner with m edges on $p \leq m$ pairs must have at least $\Omega(p/(n\log n))$ rounds, as long as $m \leq n^{1+\frac{4}{9\alpha+3\beta-10}}$, where the constant in the Ω notation depends on α,β.

We believe the difficulty in obtaining this lower bound arises from the fact that standard reductions from set-disjointness and equality are unsuitable for this task. At a high level, in most standard reductions the problem boils down to deciding the existence of an edge (which can represent, e.g., the intersecting element between the inputs); when constructing spanners, no specific edge needs to be added to the spanner or omitted from it, so the solution is allowed a considerable amount of slack that is not affected by any particular edge alone.

Instead, to obtain our lower bound, we define a new communication complexity problem that reduces to computing a sparse spanner, and prove a lower bound on its communication complexity using information theory. In this new problem, which we call $\mathsf{PART\text{-}COMP}_{m,p}$, Alice has a set $x \subseteq \{1,\ldots,m\}$ of size $|x| = p$, and Bob has to output a set $y \subseteq \{1,\ldots,m\}$ of size $|y| = m/2$ so that $x \cap y = \emptyset$. We show that any protocol that solves $\mathsf{PART\text{-}COMP}_{m,p}$ must convey $\Omega(p)$ bits of information about the set x. This technique significantly extends the current toolbox used for obtaining lower bounds for the CONGEST model. As such, we believe it may find additional applications, especially in obtaining lower bounds for computing in this model.

Roadmap: We conclude this section with a further discussion of related work. Section 2 contains the definition of the model and some basic routines. In Section 3 we present distributed algorithms for computing the various types of spanners discussed above. In Section 4 we present our new lower bounds, and we conclude with a short discussion in Section 5.

1.3 Related Work

Sparse spanners with a small multiplicative stretch are well-understood: Althöfer et al. [ADD+93] in 1993 showed that any weighted graph G on n vertices has a spanner of size $O(n^{1+1/k})$ with multiplicative stretch $2k-1$, for every integer $k \geq 1$. Since then, several works [BS07, DHZ00, EP04, Elk05, Knu14, Pet09, RTZ05, RZ11, TZ06] have considered the problem of efficiently constructing sparse spanners with small stretch and have used spanners in the applications of computing approximate distances and approximate shortest paths efficiently.

For unweighted graphs, one seeks spanners where the stretch is purely additive and as mentioned earlier, an almost tight bound of $n^{4/3}$ is known for how sparse a purely additive spanner can be. Bollobás et al. [BCE05] were the first to study a variant of pairwise preservers called *distance preservers*, where the set of relevant pairs is $\mathcal{P} = \{(u,v) : \delta_G(u,v) \geq d\}$, for a given parameter d. Coppersmith and Elkin [CE06] showed pairwise preservers of size $O(n\sqrt{|\mathcal{P}|})$ and $O(n+|\mathcal{P}|\sqrt{n})$ for any $\mathcal{P} \subseteq V \times V$. For $|P| = \omega\left(n^{3/4}\right)$, the bound of $O(n\sqrt{|\mathcal{P}|})$ for pairwise preservers has very recently been improved to $O(n^{2/3}|\mathcal{P}|^{2/3} + n|\mathcal{P}|^{1/3})$ by Bodwin and Williams [BW16].

The problem of designing sparse pairwise spanners was first considered by Cygan et al. [CGK13] who showed a tradeoff between the additive stretch and size of the spanner. The current sparsest pairwise spanner with purely additive stretch has size $O(n|\mathcal{P}|^{1/4})$ and additive stretch 6 [Kav15]. Woodruff [Woo10] and Abboud and Bodwin [AB16b, AB16a] showed lower bounds for additive spanners and pairwise spanners. Parter [Par14] showed sparse *multiplicative* sourcewise spanners and a lower bound of $\Omega(n|S|^{1/k}/k)$ on the size of a sourcewise spanner with additive stretch $2(k-1)$, for any integer $k \geq 1$.

Distributed construction of sparse spanners with multiplicative stretch was addressed in several studies [BKMP10, BS07, DG08, DGP07, DGPV08, DMP+05, Pet10]. Constructions of (α, β)-spanners were addressed in [BKMP10, DGPV09, Pet10]. Towards the goal of obtaining purely additive spanners, for which $\alpha = 1$, Elkin and Peleg [EP04] introduced nearly-additive spanners, for which $\alpha = 1 + \epsilon$. Additional distributed constructions of nearly-additive spanners are given in [Elk05, EZ06, DGPV09, Pet10]. Finally, somewhat related, are constructions of various spanners in the streaming model, and in dynamic settings, both centralized and distributed [BKS12, Bas08, BS08, Elk07a, Elk07b].

In his seminal paper, Pettie [Pet10] presents lower bounds for the number of rounds needed by distributed algorithms in order to construct several families of spanners. Specifically, it is shown that computing an all-pair additive β-spanner with size $n^{1+\rho}$ in expectation, for a constant β, requires $\Omega\left(n^{(1-\rho)/2}\right)$ rounds of communication. Because this is an indistinguishability-based lower bound, it holds even for the less restricted LOCAL mode, where message lengths can be unbounded.

The lower bound is obtained by showing an n-node graph with diameter $D = \Theta\left(n^{(1-\rho)/2}\right)$ where, roughly speaking, removing *wrong* edges induces a stretch that is too large, and identifying these wrong edges takes $\Omega(D)$ rounds. This gives a lower bound of $\min\left\{\Omega\left(n^{(1-\rho)/2}\right), \Omega\left(D\right)\right\}$ rounds. Examining the construction in detail, it is not hard to show it works for other types of spanners as well: even for a single pair of nodes, or a set S of size 2, at least $\Omega(D)$ rounds are necessary in order to avoid removing wrong edges.

2 Preliminaries

The Model: The distributed model we assume is the well-known CONGEST model [Pel00]. Such a system consists of a set of n computational units, who exchange messages according to an undirected *communication graph* $G = (V, E)$, $|V| = n$, where nodes represent the computational units and edges the communication links. Each node has a unique identifier which can be encoded using $O(\log n)$ bits. The diameter of G is denoted by D.

When the computation starts, each node knows its own identifier and the identifiers of its neighbors; when there is a set S of nodes or a set $\mathcal{P}$ of node-pairs involved in the computation, it also knows if it belongs to S, or all the pairs in $\mathcal{P}$ it belongs to. The computation proceeds in rounds, where in each round each node sends an $O(\log n)$-bits message to each of its neighbors, receives a message from each neighbor, and performs a computation. We use the number of rounds as our complexity measure, while ignoring the local computation time; however, in our algorithms all local computations take polynomial time. When the computation ends, each node knows which of its neighbors is also its neighbor in the new graph $H = (V, E')$ generated. We do not assume that the global structure of H is known to any of the nodes.

Clustering and BFS: The first building block in all of our algorithms is *clustering*. A *cluster* C_i around a *cluster center* c_i is a subset of $\Gamma_G(c_i)$, the set of neighbors of c_i in G (which does not include c_i itself). A node belonging to a cluster is *clustered*, while the other nodes of G are *unclustered*. We use $\mathcal{C}$ to denote the set of cluster centers and $\hat{\mathcal{C}}$ to denote the set of clusters.

In the *clustering* phase of our algorithms we divide some of the nodes into clusters. We create a new graph containing all the edges connecting a clustered node to its cluster center, and all the edges incident on unclustered nodes.

Another building block is *BFS trees*. A BFS tree in a graph G, rooted at a node r, consists of shortest paths from r to all other nodes in G. The process of creating a BFS tree, known as BFS search, is well-known in the sequential setting. In the distributed setting, a single BFS tree can be easily constructed by a techniques called *flooding* (see, e.g. [Pel00, §3]), and a celebrated result of Holzer and Wattenhofer [HW12] asserts that multiple BFS trees, rooted at a set S of nodes, can be constructed in $O(|S| + D)$ rounds. Here, D denotes the diameter of the graph, i.e. the maximal distance between two nodes. We use this technique to add BFS trees to the spanner we construct, and to measure distances in the original graph.

3 Building Spanners

In this section we present distributed algorithms for building several types of additive spanners. For each spanner, we first describe a template for constructing it independently of a computational model and analyze its stretch and number of edges. Then, we provide a distributed implementation of the algorithm in the CONGEST model and analyze its running time.

In a nutshell, our algorithms have three steps: first, each node tosses a coin to decide if it will serve as a cluster center; second, each cluster center tosses another coin to decide if it will serve as a root of a BFS tree; third, add to the current graph edges that are part of certain short paths. The parameters of the coins and the meaning of "short" are carefully chosen, depending on the input to the problem and the desired stretch.

Proving that the algorithms perform well is about analyzing the probability of failure. This analysis uses the graph structure as well as standard concentration bounds. In all of our algorithms, c is a constant that can be chosen according to the desired exponent of $1/n$ in the failure probability.

3.1 A $(+2)$-Sourcewise Spanner

Our first algorithm constructs a $(+2)$-sourcewise spanner. Given a set $S \subseteq V$, the algorithm returns a subgraph H of G satisfying $\delta_H(s,v) \leq \delta_G(s,v)+2$ for all $(s,v) \in S \times V$, with guarantees as given in the following theorem.

Theorem 1. *Given a graph G on n nodes and a set of sources S, a $(+2)$-sourcewise spanner with $O(n^{5/4}\,|S|^{1/4}\log^{3/4} n)$ edges can be constructed in $O(|S|+D)$ rounds in the CONGEST model w.h.p.*

This is only a factor $O(\log^{1/2} n)$ more than the number of edges given by the best sequential algorithm known for this type of spanners [KV13]. Lemmas 2 and 3 analyze the size and stretch of Algorithm 2S given below. The number of rounds of its distributed implementation is analyzed in Lemma 4, giving Theorem 1.

3.1.1 Algorithm 2S

Input: a graph $G=(V,E)$; a set of source nodes $S \subseteq V$
Output: a subgraph H
Initialization: $n=|V|$, $h=(n\,|S|)^{1/4}\log^{3/4} n$, and $H=(V,\emptyset)$

Clustering Pick each node as a *cluster center* w.p. $\frac{c\log n}{h}$, and denote the set of selected nodes by $\{c_1,c_2,\ldots\}$. For each node $v \in V$, choose a neighbor c_i of v which is a cluster center, if such a neighbor exists, add the edge (v,c_i) to H, and add v to C_i; if none of the neighbors of v is a cluster center, add to H all the edges v belongs to.

BFS Pick each cluster center as a *root* of a BFS tree w.p. $\frac{h^2}{cn\log n}$, and add to H a BFS tree rooted at each chosen root.

Path Buying For each source-cluster pair $(s,C_i) \in S \times \hat{\mathcal{C}}$: build a temporary set of paths, containing a single, arbitrary shortest path from s to each $x \in C_i$; omit from the set all paths with more than $\frac{2c^2 n\log^2 n}{h^2}$ missing edges (i.e. edges in G but not in H); if any paths are left, add to H the shortest among them.

3.1.2 Analysis of Algorithm 2S

We now study the properties of the spanner H created by the algorithm; in the next section, we describe the implementation of the different phases in the CONGEST model and analyze the running time of the algorithm.

Lemma 2. *Given a graph $G=(V,E)$ with $|V|=n$ and a set $S \subseteq V$, Algorithm 2S outputs a graph $H=(V,E')$, $E' \subseteq E$, with $|E'|=O(n^{5/4}\,|S|^{1/4}\log^{3/4} n)$ edges w.p. at least $1-O(n^{-c+2})$.*

Proof. The algorithm starts with $H = (V, \emptyset)$, and adds to it only edges from G. We analyze the number of edges added in each phase.

In the first part of the clustering phase, each node adds to H at most one edge, connecting it to a single cluster center, for a total of $O(n)$ edges. Then, the probability that a node of degree at least h is left unclustered is at most $(1 - \frac{c \log n}{h})^h$ which is $O(n^{-c})$. A union bound implies that all nodes of degree at least h are clustered w.p. $1 - O(n^{-c+2})$, and thus the total number of edges added to H by unclustered nodes in the second part of the clustering phase is $O(nh)$ w.p. $1 - O(n^{-c+2})$.

A node becomes a root in the BFS phase if it is chosen as a cluster center and then as a root, which happens with probability $\frac{c \log n}{h} \cdot \frac{h^2}{cn \log n} = \frac{h}{n}$. Letting T denote the set of trees gives $\mathbb{E}[|T|] = \frac{h}{n} \cdot n = h$, and a Chernoff bound implies that $\Pr[|T| > 4h] \leq \exp(-h)$. As $h \geq c \log n$ we have $\exp(-h) = O(n^{-c})$, and the BFS phase adds at most $4h$ trees, which are $O(nh)$ edges.

Finally, each of the n nodes is chosen as a cluster center with probability $\frac{c \log n}{h}$, so $\mathbb{E}[|\mathcal{C}|] = \frac{cn \log n}{h}$. A Chernoff bound implies $\Pr[|\mathcal{C}| > \frac{4cn \log n}{h}] \leq \exp(-\frac{cn \log n}{h})$; as $h < n^{3/4}$, we have $\exp(-\frac{cn \log n}{h}) < \exp(-cn^{1/4}) = o(n^{-c})$. For each pair in $S \times \hat{\mathcal{C}}$, at most $\frac{2c^2 n \log^2 n}{h^2}$ edges are added in the path buying phase, for a total of $O(|S| \cdot \frac{n \log n}{h} \cdot \frac{n \log^2 n}{h^2}) = O(\frac{|S| n^2 \log^3 n}{h^3})$ edges.

Substituting $h = (n\,|S|)^{1/4} \log^{3/4} n$ gives a total of $O(n^{5/4}\,|S|^{1/4} \log^{3/4} n)$ edges, as claimed. □

Lemma 3. *Given a graph $G = (V, E)$ with $|V| = n$ and a set $S \subseteq V$, the graph H constructed by Algorithm* 2S *satisfies $\delta_H(s, v) \leq \delta_G(s, v) + 2$ for each pair $(s, v) \in S \times V$ w.p. at least $1 - O(n^{-c})$.*

Proof. Consider a shortest path ρ between $s \in S$ and $v \in V$ in G.

If ρ has more than $\frac{2c^2 n \log^2 n}{h^2}$ missing edges in H after the clustering phase, then it traverses more than $\frac{c^2 n \log^2 n}{h^2}$ clusters, as otherwise there is a shorter path between s and v in G. The probability that none of the centers of these clusters is chosen as a root in the BFS phase is at most $\left(1 - \frac{h^2}{cn \log n}\right)^{c^2 n \log^2 n / h^2} = O(\exp(-c \log n)) = O(n^{-c})$. Let C_i be a cluster that ρ traverses, and let u be a node in $\rho \cap C_i$. Adding a BFS tree rooted at c_i ensures that $\delta_H(s, c_i) = \delta_G(s, c_i) \leq \delta_G(s, u) + 1$ and similarly $\delta_H(c_i, v) = \delta_G(c_i, v) \leq \delta_G(u, v) + 1$. By the triangle inequality

$$\delta_H(s, v) \leq \delta_H(s, c_i) + \delta_H(c_i, v) \leq \delta_G(s, u) + \delta_G(u, v) + 2$$

which equals $\delta_G(s, v) + 2$ since u is on ρ. This completes the proof for ρ with many missing edges.

Consider the complementary case, where ρ has at most $\frac{2c^2 n \log^2 n}{h^2}$ missing edges in H after the clustering phase. If ρ traverses no clusters, then it is contained in H, and $\delta_H(s, v) = \delta_G(s, v)$. Otherwise, if v belongs to some cluster C_i, then there is a node $v' \in C_i$ (possibly v itself) such that the shortest path between s and v' is added to the graph H in the path buying phase. The nodes v and v' both belong to the same cluster, so $\delta_H(v, v') \leq 2$ and the triangle inequality implies $\delta_H(s, v) \leq \delta_H(s, v') + \delta_H(v', v) \leq \delta_H(s, v') + 2$, as claimed. Finally, consider the case where ρ traverses at least one cluster and v is unclustered; let u be the clustered node closest to v on ρ. The sub-path from v to u is contained in H, so $\delta_H(v, u) = \delta_G(v, u)$ and by the previous analysis $\delta_H(s, u) \leq \delta_G(s, u) + 2$; the triangle inequality implies

$$\delta_H(s, v) \leq \delta_H(s, u) + \delta_H(u, v) \leq \delta_G(s, u) + 2 + \delta_G(u, v)$$

and since u is on ρ, we have $\delta_G(s, u) + \delta_G(u, v) = \delta_G(s, v)$ and the claim follows. □

3.1.3 Implementing Algorithm 2S in the CONGEST Model

We now discuss the implementation of Algorithm 2S in the CONGEST model.

Lemma 4. *Algorithm* 2S *can be implemented in* $O(|S| + D)$ *rounds in the CONGEST model, w.p. at least* $1 - o(n^{-c})$.

Proof. We present distributed implementations for each of the phases in Algorithm 2S, and analyze their running time.

Preprocessing In order to run the algorithm properly, we need each node to know the parameter h, which in turn depends on $|S|$ and n. These parameters are not given in advance to all graph nodes, but they can be gathered along a BFS tree rooted at a predetermined node, e.g. the node with minimal identifier, and then spread to all the nodes over the same tree. This is done in $O(D)$ rounds.

Clustering The clustering phase is implemented as follows: first, each node becomes a cluster center w.p. $\frac{c \log n}{h}$ and sends a message to all its neighbors; then, each node that gets at least one message joins a cluster of one of its neighbors, by sending a message to that neighbor and adding their connecting edge to the graph; finally, nodes that are not neighbors of any cluster center send a message to all their neighbors and add all their incident edges to the graph. The round complexity of this phase is constant.

BFS Each cluster center becomes a root of a BFS tree w.p. $\frac{ch^2}{n \log n}$, which is done without communication. Then, all BFS roots run BFS searches in parallel. The number of BFS trees is $O(h)$ w.p. $o\left(n^{-c}\right)$, as seen in the proof of Lemma 2, and this number of BFS searches can be run in parallel in $O(D + h)$ rounds, using an algorithm of Holzer and Wattenhofer [HW12, §6.1]. Their algorithm outputs the distances along the BFS trees, whereas we wish to mark the BFS tree edges and add them to the graph; this requires a simple change to the algorithm, which does not affect its correctness or asymptotic running time.

Path Buying This phase starts with measuring all the distances between pairs of nodes in $S \times V$, and the number of missing edges in each shortest path measured. To find all distances from a node $s \in S$ to all other nodes, we run a BFS search from s; moreover, we augment each BFS procedure with a counter that counts the missing edges in each path from the root to a node on the BFS tree. Running BFS searches from all the nodes of S is done in $O(|S| + D)$ rounds, as before, and adding a counter does not change the time complexity. When a node $v \in V$ receives a message of a BFS initiated by some $s \in S$, it learns its distance from s and the number of missing edges on one shortest path from s to v, which lies within the BFS tree; we refer to this path as *the* shortest path from s to v.

After all the BFS searches complete, each clustered node x sends to its cluster center the distance from each $s \in S$ to x, and the number of missing edges on the corresponding path. This sub-phase takes $O(|S|)$ rounds to complete.

Each cluster center c_i now knows, for each $s \in S$, the length of the shortest path from s to each $x \in C_i$, and the number of missing edges in each such path; it then locally chooses the shortest among all paths with at most $\frac{2c^2 n \log^2 n}{h^2}$ missing edges. Finally, for each chosen (x, s) path, c_i sends

a message to x containing the identifier of s. All BFS searches are now executed backwards, by sending all the messages in opposite direction and order; when x runs backwards the BFS search initiated by s, it marks the message to his parent with a "buy" bit, which is passed up the tree and makes each of its receivers add the appropriate edge to the graph. This sub-phase requires $O(|S| + D)$ rounds as well.

In total, the running time of the algorithm is $O(h + |S| + D)$, w.p. at least $1 - o(n^{-c})$, which completes the proof for the case $h = O(|S|)$. In the case $h = \Omega(|S|)$, we can replace the algorithm by a simpler algorithm that returns the union of BFS trees rooted at all nodes of S. This creates a graph that exactly preserves all distances among pairs $(s, v) \in S \times V$, and takes $O(|S| + D)$ rounds to complete. The number of edges in the created spanner is $O(n\,|S|)$, and the assumption $h = \Omega(|S|)$ implies $O(n\,|S|) = O(nh)$, as desired. □

3.2 A (+4)-All-Pairs Spanner

Recall that a subgraph H of G is a (+4)-*all-pairs* spanner if $\delta_H(u, v) \leq \delta_G(u, v) + 4$ for all pairs $(u, v) \in V \times V$. We present an algorithm, based on Algorithm `2S`, which builds a (+4)-all-pairs spanner and has the properties guaranteed by the following theorem.

Theorem 5. *Given a graph G on n nodes, a (+4)-all-pairs spanner with $O\left(n^{7/5} \log^{4/5} n\right)$ edges can be constructed in $O(n^{3/5} \log^{1/5} n + D)$ rounds in the CONGEST model w.h.p.*

The main idea is that cluster centers are now sources for a (+2)-sourcewise spanner, which, as we show, promises a (+4)-stretch to all pairs. Lemmas 6 and 7 analyze the size and the stretch of Algorithm `4AP` below. Lemma 8 analyzes the running time of its distributed implementation, completing the proof of Theorem 5.

3.2.1 Algorithm `4AP`

Input: $G = (V, E)$
Output: a subgraph H
Initialization: $n = |V|$, $h = n^{2/5} \log^{4/5} n$, and $H = (V, \emptyset)$

Clustering Run clustering as in Algorithm `2S`.

BFS and Path Buying Run the BFS and path buying phases from Algorithm `2S`, with cluster centers as sources, i.e. $S = \mathcal{C}$.

3.2.2 Analysis of Algorithm `4AP`

Lemma 6. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm `4AP` outputs a graph $H = (V, E')$, $E' \subseteq E$, with $|E'| = O\left(n^{7/5} \log^{4/5} n\right)$ edges w.p. at least $1 - O(n^{-c+2})$.*

Proof. The lemma follows from the proof of Lemma 2: in Algorithm `4AP`, S is the set of all cluster centers, whose amount is $|\mathcal{C}|$, and by the proof $|\mathcal{C}| \leq \frac{4cn \log n}{h}$ w.h.p. Substituting $|S| = O\left(\frac{n \log n}{h}\right)$ and $h = n^{2/5} \log^{4/5} n$ in Lemma 2, we get that the graph created by Algorithm `4AP` contains $O(n^{7/5} \log^{4/5} n)$ edges w.p. at least $1 - O(n^{-c+2})$. □

Lemma 7. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm* `4AP` *outputs a graph H satisfying $\delta_H(u, v) \leq \delta_G(u, v) + 4$ for each pair of vertices $(u, v) \in V \times V$ w.p. at least $1 - O(n^{-c})$.*

Proof. Let $(u, v) \in V \times V$ be an arbitrary pair of nodes, and set a shortest path ρ in G between them.

If ρ is not incident on any clustered node, then all its nodes are unclustered and all its edges are present in H. Otherwise, let x be the first clustered node on ρ, when traversing it from u to v, and let C_i be the cluster containing x. The sub-path of ρ from u to x exists in H, as all nodes on this sub-path except for x are unclustered; the distance from c_i to v satisfies $\delta_H(c_i, v) \leq \delta_G(c_i, v) + 2$, as the stretch of (c_i, v) in H is at most 2 by Lemma 3, w.p. at least $1 - O(n^{-c})$. The triangle inequality completes the proof:

$$\begin{aligned}\delta_H(u, v) &\leq \delta_H(u, x) + \delta_H(x, c_i) + \delta_H(c_i, v) \\ &\leq \delta_G(u, x) + 1 + \delta_G(c_i, v) + 2 \\ &\leq \delta_G(u, x) + \delta_G(c_i, x) + \delta_G(x, v) + 3 = \delta_G(u, v) + 4.\end{aligned}$$

□

3.2.3 Implementing Algorithm `4AP`

Running Algorithm `4AP` is done by executing Algorithm `2S` with a specific set S; thus, their running times are identical, as stated in the next lemma.

Lemma 8. *Algorithm* `4AP` *can be implemented in the CONGEST model in $O(n^{3/5} \log^{1/5} n + D)$ rounds w.p. at least $1 - o(n^{-c})$.*

Proof. The lemma follows from the proof of Lemmas 2 and 4: in Algorithm `4AP`, S is the set of cluster centers, whose amount is $|\mathcal{C}|$, and by the proof of lemma 2, $|\mathcal{C}| \leq \frac{4cn \log n}{h}$ w.h.p. Substituting $|S| = O(\frac{n \log n}{h})$ and $h = n^{2/5} \log^{4/5} n$ in Lemma 4, we get that the algorithm completes in $O(n^{3/5} \log^{1/5} n + D)$ rounds, with the claimed probability. □

3.3 A (+2)-Pairwise Spanner

Recall that a (+2)-pairwise spanner, for a set of pairs $\mathcal{P} \subseteq V \times V$, is subgraph H of G satisfying $\delta_H(u, v) \leq \delta_G(u, v) + 2$ for all pairs $(u, v) \in \mathcal{P}$. Recall that $\tau(\mathcal{P})$ denotes the number of nodes appearing in $\mathcal{P}$, i.e. $\tau(\mathcal{P}) = |\{u \mid \exists v : \{u, v\} \in \mathcal{P}\}|$.

We present a distributed algorithm, Algorithm `2P`, which returns a (+2)-pairwise spanner with the properties described in the following theorem.

Theorem 9. *Given a graph G on n nodes and a set $\mathcal{P}$ of pairs of nodes in G, a (+2)-pairwise spanner with $O(n \left|\mathcal{P}\right|^{1/3} \log^{2/3} n)$ edges can be constructed in $O(\tau(\mathcal{P}) + D)$ rounds in the CONGEST model w.h.p.*

If $\tau(\mathcal{P}) < 2c^2 \left|\mathcal{P}\right|^{1/3} \log^{2/3} n$, achieving the desired spanner is simple: for each u appearing in $\mathcal{P}$, add to H a BFS from u. The number of edges is $O\left(n\tau(\mathcal{P})\right) = O\left(n \left|\mathcal{P}\right|^{1/3} \log^{2/3} n\right)$, the stretch is 0 for all pairs in $\mathcal{P}$, and the running time is $O(\tau(\mathcal{P}) + D)$, as desired. Otherwise, Lemmas 10 and 11 prove the claimed size and stretch of Algorithm `2P` below. Lemma 12 proves the running time of its distributed implementation, giving Theorem 9.

3.3.1 Algorithm 2P

Input: $G = (V, E)$, $\mathcal{P} \subseteq V \times V$
Output: a subgraph H
Initialization: $n = |V|$, $h = |\mathcal{P}|^{1/3} \log^{2/3} n$, and $H = (V, \emptyset)$

Clustering and BFS Run clustering and add BFS trees from selected cluster centers, as in Algorithm 2S.

Path Buying For each pair $(u, v) \in \mathcal{P}$, if the shortest path between u and v in G has at most $\frac{2c^2 n \log^2 n}{h^2}$ missing edges in H, add it to H.

3.3.2 Analysis of Algorithm 2P

Lemma 10. *Given a graph $G = (V, E)$ with $|V| = n$ and a set $\mathcal{P} \subseteq V \times V$, Algorithm 2P outputs a graph $H = (V, E')$, $E' \subseteq E$, with $|E'| = O(n\,|\mathcal{P}|^{1/3} \log^{2/3} n)$ edges, w.p. at least $1 - O\left(n^{-c+2}\right)$.*

Proof. The lemma follows from the proof of Lemma 2: the clustering and BFS phases add $O(nh)$ edges to the graph w.p. at least $1 - O\left(n^{-c+2}\right)$, as long as $c \log n \leq h \leq n^{3/4}$. The first inequality comes from the comment after the statement of Theorem 9, the fact that $\tau(\mathcal{P}) \leq 2\,|\mathcal{P}|$ and the choice of h, and the second inequality is immediate.

In the path buying phase, at most $\frac{2c^2 n \log^2 n}{h^2}$ edges are added for each pair in $\mathcal{P}$, for a total of $O\left(\frac{|\mathcal{P}| n \log^2 n}{h^2}\right)$ edges. Substituting $h = |\mathcal{P}|^{1/3} \log^{2/3} n$, we get a total of $O\left(n\,|\mathcal{P}|^{1/3} \log^{2/3} n\right)$ edges in H. □

Lemma 11. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm 2P outputs a graph H satisfying $\delta_H(u, v) \leq \delta_G(u, v) + 2$ for each pair of vertices $(u, v) \in \mathcal{P}$, w.p. at least $1 - O\left(n^{-c}\right)$.*

Proof. Let $(u, v) \in \mathcal{P}$ be an arbitrary pair of nodes, and fix a shortest path ρ in G between them.

If ρ has at most $\frac{2c^2 n \log^2 n}{h^2}$ missing edges in H before the path buying phase, it is added to H, and $\delta_G(u, v) = \delta_H(u, v)$. Otherwise, ρ has more than $\frac{2c^2 n \log^2 n}{h^2}$ missing edges before the BFS phase, so it traverse at least $\frac{c^2 n \log^2 n}{h^2}$ clusters. As in the proof of Lemma 3, at least one of the corresponding cluster centers is chosen as a root of a BFS tree w.p. at least $1 - O\left(n^{-c}\right)$, and $\delta_H(u, v) \leq \delta_G(u, v) + 2$, as claimed. □

3.3.3 Implementing Algorithm 2P

Lemma 12. *Algorithm 2P can be implemented in $O(\tau(\mathcal{P}) + |\mathcal{P}|^{1/3} \log^{2/3} n + D)$ rounds in the CONGEST model w.p. at least $1 - o\left(n^{-c}\right)$.*

Proof. We can implement the clustering and path buying phases in $O(h + D) = O(\tau(\mathcal{P}) + D)$ rounds with success probability $1 - o\left(n^{-c}\right)$, as seen in the proof of Lemma 4. In order to count missing edges in paths, we run a BFS search in G from each node appearing in $\mathcal{P}$. Then, the BFS search is run backwards, and is used to add the "cheap" paths: for a pair (u, v) in $\mathcal{P}$, if the BFS from v arrives at u traversing at most $\frac{2c^2 n \log^2 n}{h^2}$ missing edges, then u sends back a "buy" message up the tree, and the path is added. We may end up adding two shortest path for a pair $(u, v) \in \mathcal{P}$,

but this does not affect the asymptotic number of edges or the time complexity. This phase is implemented in $O(\tau(\mathcal{P}) + D)$ rounds, by running the $\tau(\mathcal{P})$ BFS searches in parallel. □

3.4 A (+4)-Pairwise Spanner

We present an algorithm for constructing a (+4)-pairwise spanner, with the parameters described by the following theorem.

Theorem 13. *Given a graph G on n nodes and a set $\mathcal{P}$ of pairs, a (+4)-pairwise spanner with $O\left(n\left|\mathcal{P}\right|^{2/7}\log^{6/7} n\right)$ edges can be constructed in $O(\tau(\mathcal{P}) + D)$ rounds in the CONGEST model w.h.p.*

If $|\mathcal{P}| < \log^4 n$, the (+2)-pairwise spanner from Theorem 9 is sparser than the one promised by Theorem 13, and can be constructed in the same running time. Otherwise, Lemmas 14 and 15 show the claimed size and stretch of Algorithm `4P` below, which together with Lemma 16, which analyzes the running time of its distributed implementation, proves Theorem 13.

3.4.1 Algorithm `4P`

Input: a graph $G = (V, E)$; a set of pairs $\mathcal{P} \subseteq V \times V$
Output: a subgraph H
Initialization: $n = |V|$, $h = \left|\mathcal{P}\right|^{2/7}\log^{6/7} n$, $\ell = \frac{n\log^3 n}{h^{5/2}}$ and $H = (V, \emptyset)$

Clustering and BFS Run clustering and add BFS trees from selected cluster centers, as in Algorithm `2S`.

Prefix-Suffix Buying For each pair $(u, v) \in \mathcal{P}$, let ρ be a shortest path from u to v. Add to H the first ℓ missing edges and the last ℓ missing edges in ρ.

Choosing Cluster Centers Construct a set A of cluster centers by adding to it each cluster center independently w.p. $\frac{16c\log n}{\ell}$.

Path Buying For each pair $(c_1, c_2) \in A \times A$: fix a set of paths containing a single shortest path from c_1 to each $x \in C_2$; omit all paths with more than $\frac{2c^2 n\log^2 n}{h^2}$ missing edges in H; if any paths are left, add to H the shortest among them.

3.4.2 Analysis of Algorithm `4P`

Lemma 14. *Given a graph $G = (V, E)$ with $|V| = n$ and a set $\mathcal{P} \subseteq V \times V$, Algorithm `4P` outputs a graph $H = (V, E')$, $E' \subseteq E$, with $|E'| = O\left(n\left|\mathcal{P}\right|^{2/7}\log^{6/7} n\right)$ edges w.p. at least $1 - O\left(n^{-c+2}\right)$.*

Proof. The clustering and BFS phases add $O(nh)$ edges to the graph w.p. at least $1 - O\left(n^{-c+2}\right)$, as seen in the proof of Lemma 2, as long as $c\log n < h < n^{3/4}$. The first inequality comes from the discussion below the statement of Theorem 13, and the second is immediate.

In the prefix-suffix buying phase, at most $O(\ell)$ edges are bought for each pair in $\mathcal{P}$, for a total of $O\left(|\mathcal{P}| \cdot \frac{n\log^3 n}{h^{5/2}}\right) = O\left(n\left|\mathcal{P}\right|^{2/7}\log^{6/7} n\right)$ edges.

Finally, in the path buying phase we add to H at most $|A|^2$ paths, with $O\left(\frac{n\log^2 n}{h^2}\right)$ missing edges in each. Each node is chosen to be a cluster center w.p. $\frac{c\log n}{h}$, and then to enter A w.p. $\frac{16c\log n}{\ell}$, so $\mathbb{E}[|A|] = n\cdot\frac{c\log n}{h}\cdot\frac{16c\log n}{\ell} = \frac{16c^2 n\log^2 n}{h\ell} = 16c^2\left|\mathcal{P}\right|^{3/7}\log^{2/7} n$. A Chernoff bound implies $\Pr\left[|A| > 64c^2\left|\mathcal{P}\right|^{3/7}\log^{2/7} n\right] \leq \exp\left(-16c^2\left|\mathcal{P}\right|^{3/7}\log^{2/7} n\right) = O(n^{-c})$; the last equality holds under the assumption $|\mathcal{P}| \geq \log n$, as discussed below the statement of Theorem 9. Hence, the number of edges added in the path buying step is $O\left(\left(\left|\mathcal{P}\right|^{3/7}\log^{2/7} n\right)^2\cdot\frac{n\log^2 n}{h^2}\right) = O(\left(n\left|\mathcal{P}\right|^{2/7}\log^{6/7} n\right)$ edges. In total, H has $O(\left(n\left|\mathcal{P}\right|^{2/7}\log^{6/7} n\right)$ edges w.p. at least $1 - O\left(n^{-c+2}\right)$. □

Lemma 15. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm* `4P` *outputs a graph H satisfying $\delta_H(u,v) \leq \delta_G(u,v) + 4$ for each pair of vertices $(u,v) \in \mathcal{P}$ w.p. at least $1 - O\left(n^{-c+2}\right)$.*

Proof. Let $(u,v) \in \mathcal{P}$ be an arbitrary pair of nodes, and let ρ be an arbitrary shortest path from u to v. If ρ has at most 2ℓ missing edges in H after the clustering phase, it is added to H in the prefix-suffix buying phase and $\delta_G(u,v) = \delta_H(u,v)$.

Otherwise, the prefix of ρ with ℓ missing edges is incident on at least $\ell/2$ clusters. Each cluster center is added to A independently w.p. $\frac{16c\log n}{\ell}$ so the expected number of clusters in A which are also incident on the prefix is $8c\log n$, and a Chernoff bound implies that the probability that less than $4c\log n$ of the centers of these clusters are chosen to A is at most $O(\exp(-c\log n)) = O\left(n^{-c}\right)$. The same argument shows that the suffix of ρ is incident on a cluster in A, and a union bound implies that all prefixes and suffixes are incident on clusters in A w.p. at least $1 - O\left(n^{-c+2}\right)$.

Let c_1 be a center of a cluster in A which is incident on the prefix of ρ, and c_2 a center of a cluster incident on the suffix of ρ. Let u' and v' be nodes in $\rho \cap C_1$ and $\rho \cap C_2$ respectively, and let σ be a path between c_1 and v' in G.

If the number of edges of σ missing in H after the clustering phase is more than $\frac{2c^2 n\log^2 n}{h^2}$, then σ is incident on at least $\frac{c^2 n\log^2 n}{h^2}$ clusters. In this case, a cluster incident on σ is a source of a BFS tree w.p. at least $1 - O\left(n^{-c}\right)$, as seen in the proof of Lemma 3. Let C_i be such a cluster, then after adding the BFS trees it holds that $\delta_H(u',v') \leq 1 + \delta_H(c_1,v') \leq 1 + \delta_H(c_1,c_i) + \delta_H(c_i,v') \leq \delta_G(c_1,v')+3 \leq \delta_G(u',v')+4$, which implies $\delta_H(u,v) \leq \delta_H(u,u')+\delta_H(u',v')+\delta_H(v',v) \leq \delta_G(u,u')+\delta_G(u',v')+4+\delta_G(v',v) = \delta_G(u,v)+4$.

If σ has less than $\frac{2c^2 n\log^2 n}{h^2}$ missing edges then a path between c_1 and some $x \in C_2$ is added to H in the path buying phase, satisfying $\delta_H(c_1,x) \leq \delta_G(c_1,v') \leq \delta_G(u',v')+1$. Hence, $\delta_H(u,v) \leq \delta_H(u,u')+\delta_H(u',c_1)+\delta_H(c_1,x)+\delta_H(x,v')+\delta_H(v',v) \leq \delta_G(u,u')+1+\delta_G(u',v')+1+2+\delta_G(v',v) = \delta_G(u,v)+4$, as required. □

3.4.3 Implementing Algorithm `4P`

Lemma 16. *Algorithm* `4P` *can be implemented in $O\left(\tau(\mathcal{P}) + D\right)$ rounds in the CONGEST model w.p. at least $1 - o\left(n^{-c}\right)$.*

Proof. We can implement the clustering and path buying phases in $O(h+D)$ rounds with success probability $1 - o\left(n^{-c}\right)$, as seen in the proof of Lemma 4.

The prefix-suffix buying phase is implemented as follows: run a BFS from each u appearing in $\mathcal{P}$, counting missing edges on each path. Roll back the BFS, and when passing a node v such that $(u,v) \in \mathcal{P}$, it sends a "buy-suffix" message with a counter initiated to ℓ; when a node receives such

a counter it adds the edge to the parent to H and decreases the counter by 1, unless the edge is already on H; if the counter reaches another node in v' satisfying $(u, v') \in \mathcal{P}$, it is set to ℓ again. When the counter is 0, no more edges are bought but the counter is passed up the tree, until it arrives at a node whose count of missing edges in the original BFS was ℓ. This node replaces it with a simple "buy" message, adds the edge to its parent to H and sends it to the parent; each receiver of the "buy" message does the same, all the way to the tree root. This phase takes $O(\tau(\mathcal{P}) + D)$ rounds.

The choice of A is simple, requiring no communication.

The path buying phase is implemented similarly to its implementation in Algorithm `2S`, in the proof of Lemma 4. Measuring distances and counting missing edges on a path from each $c_1 \in A$ to each other node is done by running a BFS from each cluster center with the appropriate counters; each clustered node reports its cluster's center the above parameters in $O(|A|)$ rounds; each $c_2 \in A$ then chooses which paths to buy, and reports the appropriate node in its cluster; running the BFS searches backwards, each node may send "buy" messages up the trees, as decided by c_2. This phase takes $O(|A| + D)$ rounds.

The proof of Lemma 14 implies $|A| = O\left(|\mathcal{P}|^{3/7} \log^{2/7} n\right)$ w.p. at least $1 - O\left(n^{-c}\right)$. This, together with the assumption $|\mathcal{P}| \geq \log^4(n)$ and with the choice of h, imply $O\left(h + |A| + \tau(\mathcal{P}) + D\right) = O\left(\tau(\mathcal{P}) + D\right)$; hence, the above implementation takes $O\left(\tau(\mathcal{P}) + D\right)$ rounds with the same probability. □

3.5 Subsetwise Spanners

Recall that a $(+\beta)$-subsetwise spanner for a set $S \subseteq V$ is a subgraph H of G satisfying $\delta_H(u, v) \leq \delta_G(u, v) + 2$ for all pairs $(u, v) \in S \times S$. We show how to build such spanners for $\beta = 2$ and $\beta = 4$, with $O(n\,|S|^{2/3} \log^{2/3} n)$ edges and $O(n\,|S|^{4/7} \log^{6/7} n)$ edges respectively, in $O(|S| + D)$ rounds w.h.p.

The algorithms follow immediately from Algorithm `2P` and Algorithm `4P`: set $\mathcal{P} = S \times S$ and run Algorithm `2P` or Algorithm `4P`. The number of edges follows from the fact $|\mathcal{P}| = |S|^2$, and the running time from $\tau(\mathcal{P}) = |S|$.

Finally, note that in the case $|S| > n^{3/5} \log^{1/5} n$, Algorithm `2S` gives a sparser spanner than Algorithm `2P` in the same running time, and with a stretch of $+2$ for all $S \times V$ pairs. Similarly, when $|S| > n^{7/10} \log^{-1/10} n$, Algorithm `4AP` gives a sparser spanner than Algorithm `4P` in a shorter running time, with a stretch of $+4$ on all pairs of nodes in the graph.

Corollary 17. *Given a graph G on n nodes and a set S of nodes, a $(+2)$-subsetwise spanner with $O\left(\min\left\{n\,|S|^{2/3} \log^{2/3} n, n^{5/4}\,|S|^{1/4} \log^{3/4} n\right\}\right)$ edges and a $(+4)$-subsetwise spanner with $O\left(\min\left\{n\,|S|^{4/7} \log^{6/7} n, n^{7/5} \log^{4/5} n\right\}\right)$ edges, can both be constructed in at most $O(|S|+D)$ rounds in the CONGEST model w.h.p.*

3.6 A $(+8)$-All-Pairs Spanner

Recall that a subgraph H of G is a $(+8)$-*all-pairs* spanner if $\delta_H(u, v) \leq \delta_G(u, v) + 8$ for all pairs $(u, v) \in V \times V$. We present an algorithm, based on Algorithm `4P`, which builds a $(+8)$-all-pairs spanner and has the properties guaranteed by the next theorem.

Theorem 18. *Given a graph G on n nodes, a $(+8)$-all-pairs spanner with $O\left(n^{15/11}\log^{10/11} n\right)$ edges can be constructed in $O(n^{7/11}\log^{1/11} n + D)$ rounds in the CONGEST model w.h.p.*

Lemmas 19 and 20 provide the required size and stretch of Algorithm `8AP` below, while Lemma 21 gives the running time f its distributed implementation, proving Theorem 18.

3.6.1 Algorithm `8AP`

Input: $G = (V, E)$
Output: a subgraph H
Initialization: $n = |V|$, $h = n^{4/11}\log^{10/11} n$, and $H = (V, \emptyset)$

Clustering Run clustering as in Algorithm `4P`.

Rest of Algorithm `4P` Run the rest of Algorithm `4P` on all pairs of cluster centers, i.e. $\mathcal{P} = \mathcal{C} \times \mathcal{C}$.

3.6.2 Analysis of Algorithm `8AP`

Lemma 19. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm `8AP` outputs a graph $H = (V, E')$, $E' \subseteq E$, with $|E'| = O\left(n^{15/11}\log^{10/11} n\right)$ edges w.p. at least $1 - O(n^{-c+2})$.*

Proof. We follow the outline of previous proofs. By the proof of Lemma 2, the BFS phase adds $O(nh)$ edges, and $|\mathcal{C}| \leq \frac{4cn\log n}{h}$, w.p. at least $1 - O\left(n^{-c+2}\right)$. In Algorithm `8AP`, $|\mathcal{P}| = |\mathcal{C}|^2$, and substituting $|\mathcal{P}| = O\left(\frac{n^2\log^2 n}{h^2}\right)$ and $h = n^{4/11}\log^{10/11} n$ in Lemma 14 gives that the graph created by Algorithm `8AP` contains $O(n^{15/11}\log^{10/11} n)$ edges w.p. at least $1 - O\left(n^{-c+2}\right)$. □

Lemma 20. *Given a graph $G = (V, E)$ with $|V| = n$, Algorithm `8AP` outputs a graph H satisfying $\delta_H(u, v) \leq \delta_G(u, v) + 8$ for each pair of vertices $(u, v) \in V \times V$ w.p. at least $1 - O(n^{-c})$.*

Proof. Let $(u, v) \in V \times V$ be an arbitrary pair of nodes, and set a shortest path ρ in G between them.

If ρ is not incident on any clustered node, then all its nodes are unclustered and all its edges are present in H. Otherwise, let x be the first clustered node on ρ, when traversing it from u to v, and let y the last clustered node on ρ. Let C_1 be the cluster containing x, and C_2 the cluster containing y. The sub-paths of ρ from u to x and from y to v exist in H, as all the nodes on these sub-path except for x and y are unclustered. By Lemma 15, $\delta_H(c_1, c_2) \leq \delta_G(c_1, c_2) + 4$ as $c_1, c_2 \in \mathcal{C} = S$; moreover, $\delta_G(c_1, c_2) \leq \delta_G(x, y) + 2$ by the triangle inequality. Finally,

$$\begin{aligned}\delta_H(u, v) &\leq \delta_H(u, x) + \delta_H(x, c_1) + \delta_H(c_1, c_2) + \delta_H(c_2, y) + \delta_H(y, v)\\ &\leq \delta_G(u, x) + 1\delta_G(x, y) + 6 + \delta_G(y, v) + 1 = \delta_G(u, v) + 8\end{aligned}$$

as desired. □

3.6.3 Implementing Algorithm 8AP

Running Algorithm 8AP is done by executing Algorithm 4P with a specific set $\mathcal{P}$; thus, their running times are identical, as stated in the next lemma.

Lemma 21. *Algorithm* 8AP *can be implemented in the CONGEST model in* $O(n^{7/11}\log^{1/11} n + D)$ *rounds w.p. at least* $1 - o(n^{-c})$.

Proof. The lemma follows from the proof of Lemmas 2 and 16: in Algorithm 8AP, $\tau(\mathcal{P}) = |\mathcal{C}|$, and by the proof of Lemma 2, $|\mathcal{C}| = O\left(\frac{n\log n}{h}\right)$ w.p. at least $1 - O(n^{-c+2})$. Substituting this and $h = n^{4/11}\log^{10/11} n$ in Lemma 16, we get that the algorithm completes in $O(n^{7/11}\log^{1/11} n + D)$ rounds with the desired probability. □

4 Lower Bounds

In this section we prove lower bounds on the number of rounds that are needed for constructing spanners in the CONGEST model. All previous lower bounds for the distributed construction of spanners [Pet10] use an indistinguishability argument: while many edges should be omitted from the graph in order to create a sparse spanner, there are few edges that must not be omitted. However, in order to distinguish these few edges from the rest, some nodes must learn a considerable part of the graph. In a nutshell, the heart of the proof is that information must travel a constant portion of the diameter D, and thus the lower bound is $\Omega(D)$.

The lower bounds from [Pet10] apply also to the LOCAL model, where the message sizes are unbounded. Here, we present the first lower bound that is specific for the CONGEST model. As in previous lower bounds for the CONGEST model, our proof uses a reduction from a communication complexity problem. However, previous lower bounds used reductions either from the equality problem [PR99] or from set-disjointness, e.g., [DSHK+12, DKO14, FHW12, HP14, CGK14, GK13]. These seem unsuitable for our purposes, and hence we diverge from this approach and define a new communication complexity problem we call *partial complement*. We bound the communication complexity of this problem from below, using information theory.

We first prove a lower bound for the construction of a $(+2)$-pairwise spanner. Then, we generalize the bound for the construction of an (α, β)-pairwise spanner, for any $\alpha \geq 1, \beta \geq 0$.

4.1 A Communication Complexity Problem

Let m, p be two positive integers so that $p \leq m/3$. The *partial complement* communication problem, denoted $\mathsf{PART\text{-}COMP}_{m,p}$, is defined as follows: Alice has a set $x \subseteq \{1, \ldots, m\}$ of size $|x| = p$, and Bob has to output a set $y \subseteq \{1, \ldots, m\}$ of size $|y| = m/2$ so that $x \cap y = \emptyset$. Note that the goal of this communication problem is to compute a relation, not a function. In this section we prove that the randomized communication complexity of the partial complement problem is high (for formal definitions in communication complexity see the textbook [KN97]).

Theorem 22. *If* π *is a* $(1/3)$*-error randomized protocol computing* $\mathsf{PART\text{-}COMP}_{m,p}$ *then the length of* π *is at least* $p/100$.

The proof uses information theory. The basic idea is to prove that any protocol that solves $\mathsf{PART\text{-}COMP}_{m,p}$ must contain $\Omega(p)$ bits of information about the set x. We now define the basic notions we use in the proof (see the textbook [CT06] for more background and useful properties).

The *entropy* of a random variable X is defined as

$$\mathrm{H}(X) = \sum_x \Pr[X = x] \log(1/\Pr[X = x]).$$

It is well-known that entropy is maximal for the uniform distribution: if X is an element chosen at random from a set T of size t, then, by convexity,

$$\mathrm{H}[X] = \sum_{x \in T} \Pr[X = x] \log(1/\Pr[X = x]) \leq \log t. \tag{1}$$

Shannon's coding theorem says that $\mathrm{H}[X]$ is a lower bound on the expected length of a prefix-free encoding of X. The entropy of a random variable X conditioned on a random variable Y is defined as

$$\mathrm{H}(X|Y) = \mathrm{H}(X, Y) - \mathrm{H}(Y).$$

The *mutual information* between X and Y is

$$\mathrm{I}[X; Y] = \mathrm{H}[X] - \mathrm{H}[X|Y].$$

Intuitively, it measures the reduction in uncertainty about X given the value of Y.

Proof of Theorem 22. Let π be a randomized protocol computing $\mathsf{PART\text{-}COMP}_{m,p}$ with c bits of communication and error probability $1/3$. Note that we can assume without loss of generality that only Alice speaks in π. Below, we feed into π a random set X, so we may assume that π is deterministic. Denote by $\Pi = \Pi(x)$ the transcript of π with input x. Denote by $Y = Y(x)$ the subset of $\{1, \ldots, m\}$ of size $m/2$ that Bob outputs when receiving message $\Pi(x)$.

The distribution on X is defined as follows. Let $I_1, \ldots, I_p$ be p consecutive intervals in $\{1, \ldots, m\}$, each of size $k = \lfloor m/p \rfloor$. From each interval I_i, choose one element X_i uniformly at random and independently. Set

$$X = \{X_1, \ldots, X_p\}.$$

First, we compute the entropy of X. For each i, we have $\mathrm{H}[X_i] = \log k$. Thus, independence implies that

$$\mathrm{H}[X] = \sum_{i=1}^{p} \mathrm{H}[X_i] = p \log k.$$

Second, we analyze the entropy of X conditioned on Y and on the protocol outputting the correct answer. Let E be the indicator of the event

$$\{x \subset \{1, \ldots, m\} : |x| = p,\ Y(x) \cap x = \emptyset\}.$$

Note that since the protocol has an error probability bounded by $1/3$, it holds that $\Pr[E = 1] \geq 2/3$. Given y of size $m/2$, split the intervals into two sets:

$$S_1(y) = \left\{ i \in \{1, \ldots, p\} : |I_i \cap y| \leq \frac{k}{10} \right\},\ S_2(y) = \left\{ i \in \{1, \ldots, p\} : |I_i \cap y| > \frac{k}{10} \right\}.$$

Since $|S_1(y)| + |S_2(y)| = p$, it holds that

$$\begin{aligned}\frac{m}{2} &= |y| \\ &\leq |S_1(y)|\,\frac{k}{10} + |S_2(y)|\,k + k \\ &\leq \frac{pk}{10} + |S_2(y)|\,\frac{9k}{10} + k,\end{aligned}$$

which by a simple calculation implies

$$|S_2(y)| \geq \frac{4}{9}p - \frac{10}{9}.$$

In addition, for each y and $i \in S_1(y)$, using (1), we have

$$\mathrm{H}[X_i|Y = y, E = 1] \leq \log k.$$

For each y and $i \in S_2(y)$, conditioned on $E = 1$ and $Y = y$, the element X_i is chosen from $I_i \setminus y$, whose size is at most $k - \frac{k}{10} = \frac{9k}{10}$. Thus, for each y and $i \in S_2(y)$,

$$\mathrm{H}[X_i|Y = y, E = 1] \leq \log \frac{9k}{10}.$$

We can conclude this part as follows:[1]

$$\begin{aligned}\mathrm{H}[X|Y, E = 1] &\leq \sum_y \Pr[Y = y|E = 1] \left(\sum_{i \in S_1(y)} \mathrm{H}[X_i|Y = y, E = 1] + \sum_{i \in S_2(y)} \mathrm{H}[X_i|Y = y, E = 1] \right) \\ &\qquad\qquad \text{(sub-additivity)} \\ &\leq \sum_y \Pr[Y = y|E = 1] \left((p - |S_2(y)|) \log k + |S_2(y)| \log \frac{k}{10} \right) \\ &\leq \sum_y \Pr[Y = y|E = 1] \left(p \log k - |S_2(y)| \log \frac{10}{9} \right) \\ &\leq p \left(\log k - \frac{4}{9} \log \frac{10}{9} \right) + 2.\end{aligned}$$

Thus, using (1),

$$\begin{aligned}\mathrm{H}[X|Y, E] &= \Pr[E = 1] \cdot \mathrm{H}[X|Y, E = 1] + \Pr[E = 0] \cdot \mathrm{H}[X|Y, E = 0] \\ &\leq \frac{2}{3} \left(p \left(\log k - \frac{4}{9} \log \frac{10}{9} \right) + 2 \right) + \frac{1}{3} \cdot p \log k \qquad (\Pr[E = 1] \geq 2/3) \\ &\leq 2 + p \log k - \frac{4p}{9} \log \frac{10}{9}.\end{aligned}$$

We now have

$$\begin{aligned}\mathrm{I}[X; Y] &= \mathrm{I}[X; Y, E] - \mathrm{I}[E; X|Y] \qquad \text{(chain rule)} \\ &= \mathrm{H}[X] - \mathrm{H}[X|Y, E] - \mathrm{I}[E; X|Y] \\ &\geq \frac{4p}{9} \log \frac{10}{9} - 3,\end{aligned}$$

[1] Note the difference between conditioning on an event and conditioning on a random variable.

where the last inequality holds since E contains at most one bit of information. Therefore,

$$\begin{aligned} c &\geq \mathrm{H}(\Pi) && \text{(Shannon's coding theorem)} \\ &\geq \mathrm{I}(X;\Pi) \\ &\geq \mathrm{I}(X;Y) && \text{(information processing inequality)} \\ &\geq \frac{4p}{9}\log\frac{10}{9} - 3 \\ &\geq p/100; && \text{(we may assume } p > 100\text{)} \end{aligned}$$

where c is the length of π, completing the proof. □

4.2 A Lower Bound for constructing a $(+2)$-Pairwise Spanner

In this section we prove the following lower bound.

Theorem 23. *There is a constant $c > 0$ so that the following holds. Any distributed protocol for the CONGEST model with success probability at least $2/3$ which, given a graph with n nodes and a set of $p \leq cn^{3/2}$ pairs of nodes, outputs a $(+2)$-pairwise spanner with at most $cn^{3/2}$ edges, must take $\Omega(\frac{p}{n\log n})$ rounds to complete. The lower bound holds even for graphs with constant diameter.*

The theorem implies a lower bound of $\Omega\left(\sqrt{n}/\log n\right)$ on the number of rounds needed for an algorithm in the CONGEST model to output a $(+2)$-pairwise spanner, when $|\mathcal{P}| = \Theta(n^{3/2})$. For comparison, the time for constructing such a spanner using Algorithm 2P can (roughly) vary between $n^{3/4}$ and n, depending on the structure of $\mathcal{P}$.

The graph G for which the lower bound is proved is defined as follows. Let n be such that there is a finite projective plane with $n/4$ points and $n/4$ lines. Let G' be the point-line incidence graph with $n/2$ nodes (see e.g. [Mat02, §4.5]). The graph G' has $\Theta(n^{3/2})$ edges, girth 6 and diameter 3.[2] We denote the nodes of G' by $V_B = \{v'_1, \ldots, v'_{n/2}\}$. The graph G consists of G', an additional $n/2$ nodes denoted $V_A = \{v_1, \ldots, v_{n/2}\}$, and an additional $n/2$ edges of the form (v_i, v'_i).

In the pairwise spanner we construct, we wish to approximately preserve distances between pairs of nodes in V_A, i.e. $\mathcal{P} \subseteq V_A \times V_A$. The main observation is that, since the girth of G' is 6, if $e' = \{v'_i, v'_j\}$ is an edge of G' then the following holds. If $(v_i, v_j) \in \mathcal{P}$ then any $(+2)$-pairwise spanner must contain the edge e', as otherwise the distance is stretched from 3 to 7, which exceeds the required $+2$ stretch. On the other hand, if $(v_i, v_j) \notin \mathcal{P}$ then the edge e' can be safely omitted from the spanner.

Proof of Theorem 23. Fix a distributed protocol σ for constructing a $(+2)$-pairwise spanner with at most $m/2$ edges. Let G be the graph described above, and denote the edges of G' by $e_1, \ldots, e_m$.

We describe a reduction from $\mathsf{PART\text{-}COMP}_{m,p}$ to σ. Assume Alice has a set $x \subseteq \{1, \ldots, m\}$ of size p, and Bob has to output a set $y \subseteq \{1, \ldots, m\}$ of size $m/2$ satisfying $x \cap y = \emptyset$. Alice and Bob simulate σ on the graph G with the set of pairs

$$\mathcal{P} = \left\{(v_i, v_j) : \exists k \in x \; e_k = \{v'_i, v'_j\}\right\}.$$

That is, a pair (v_i, v_j) is in $\mathcal{P}$ if the corresponding pair (v'_i, v'_j) is an edge e_k whose index k is in x. Alice simulates the nodes in V_A, and Bob simulates the nodes V_B and the edges among them.

[2] The *girth* of a graph is the length of the shortest simple cycle in it.

To simulate communication on edges of the form (v_i, v'_i), Alice and Bob communicate. Note that $\mathcal{P}$ contains only pairs of nodes that are simulated by Alice.

The spanner constructed is a subgraph H of G with at most $m/2$ edges, satisfying $\delta_H(v_i, v_j) \leq \delta_G(v_i, v_j) + 2$ for all $(v_i, v_j) \in \mathcal{P}$. For each such pair, by definition of $\mathcal{P}$, we have $\delta_G(v_i, v_j) = 3$, which implies $\delta_H(v_i, v_j) \leq 5$ and $\delta_H(v'_i, v'_j) \leq 3$. The fact that G' has girth 6 implies that the edge $\{v'_i, v'_j\}$ must be in H. Let

$$y = \{k : e_k \in E_G \setminus E_H\}.$$

The spanner size implies $|y| \geq m/2$, while the above discussion implies $x \cap y = \emptyset$. Thus, Bob can output a subset of y of size $m/2$, solving the communication complexity problem.

By the communication complexity lower bound, Alice and Bob must communicate $\Omega(p)$ bits during the simulation. The number of edges they simulate together is $n/2$, and $O(\log n)$ bits are sent over each edge at each round. Thus, the protocol must take $\Omega\left(\frac{|\mathcal{P}|}{n \log n}\right)$ rounds to complete. □

4.3 Generalization: A Lower bound for constructing an (α, β)-Pairwise Spanner

Recall that an (α, β)-pairwise spanner for a graph G and a set $\mathcal{P}$ of pairs of nodes is a subgraph H of G satisfying $\delta_H(u, v) \leq \alpha\delta_G(u, v) + \beta$ for every $(u, v) \in \mathcal{P}$.

To obtain our lower bound for any $\alpha \geq 1, \beta \geq 0$, we first study the tradeoff between the girth and number of edges in a graph. The most relevant claim for this question is Erdős' girth conjecture:

Conjecture 24 (Erdős' Girth Conjecture [Erd64])**.** *For every g there is a constant c such that there exists a graph on n nodes with girth g and $cn^{1+\frac{1}{\lceil g/2 \rceil - 1}}$ edges.*

For example, for $g = 3$, the complete graph on n nodes has roughly $n^2/2$ edges, and for $g = 4$ the full bipartite graph has $n^2/4$ edges. For $g = 5$ and $g = 6$ there exist a graph with $n^{3/2}$ edges, which we used in the last section.

The conjecture is known to be true for a few values of g, while for the other values there are constructions with slightly less edges:

Theorem 25 (see, e.g. [Mat02, §15.3])**.** *For every $g \geq 3$ and $n \geq 2$ there is a graph on n vertices with girth g and $\Omega\left(n^{1+\frac{4}{3g-10}}\right)$ edges. For $g \in \{3, 4, 5, 6, 9, 10\}$ there is a constant c such that for every $n \geq 2$ there is a graph on n vertices with girth g and $cn^{1+\frac{1}{\lceil g/2 \rceil - 1}}$ edges.*

Theorem 23 and its proof extend to (α, β)-pairwise spanners for any $\alpha \geq 1$ and $\beta \geq 0$, with the appropriate choice of G'. The only thing left is to make sure the diameter is constant even if the diameter of G' is not. Bounding the diameter also allows us to derive a lower bound conditioned on Erdős' girth conjecture, a conjecture which claims nothing about the diameter of the graph.

Theorem 26. *Let $\alpha \geq 1, \beta \geq 0$ be constants, and $g = 3\alpha + \beta$. There is a constant $c = c(g)$ so that the following holds. Any distributed protocol for the CONGEST model with success probability at least $2/3$ which, given a graph with n nodes and a set of $p \leq cn^{1+\frac{4}{3g-10}}$ pairs of nodes, outputs an (α, β)-pairwise spanner with at most $cn^{1+\frac{4}{3g-10}}$ edges, must take $\Omega\left(\frac{p}{n \log n}\right)$ rounds to complete. The lower bound holds even for graphs with diameter $O(g)$.*

For $g \in \{3, 4, 5, 6, 9, 10\}$, and for any constant g if Erdős' girth conjecture is true, the bound on p and on the number of edges can be replaced by $cn^{1+\frac{1}{\lceil g/2 \rceil - 1}}$.

The theorem implies a lower bound of $\Omega\left(n/\log(n)\right)$ rounds for any algorithm in the CONGEST model which outputs pairwise preserver with $o(n^2)$ edges, when $|\mathcal{P}| = \Theta(n^2)$. The trivial algorithm that builds a BFS tree from any node appearing in $\mathcal{P}$ runs in $O(n)$ rounds and returns a spanner with $O(n^2)$ edges. Hence, we cannot expect to asymptotically improve upon the running time and upon the number of edges simultaneously.

For $(+4)$-pairwise spanners, the theorem implies a lower bound of $\Omega(n^{4/11}/\log n)$ rounds for $|\mathcal{P}| = \Theta(n^{15/11})$. Algorithm `4P` constructs such a spanner in roughly $O(n^{107/77}\log^{6/7} n)$ rounds. Assuming Erdős' conjecture, we can choose $\mathcal{P}$ satisfying $|\mathcal{P}| = \Theta(n^{4/3})$, and get a lower bound of $\Omega(n^{1/3}/\log n)$ rounds and an algorithm running in $O(n^{29/21}\log^{6/7} n)$ rounds.

The graph G for which the lower bound is proved is defined similarly to the graph in the proof of Theorem 23, with an extra construction that ensures that the diameter of G is constant. Let G' be a graph on n' nodes with girth $g = 3\alpha + \beta$ and m edges, where $m = m(n', g)$ is the maximal possible number of edges for these n' and g. Let $v'_1, \ldots, v'_{n'}$ be the nodes of G'. Add to G' a new node u and connect each node v'_i to u by a disjoint path of $\lfloor g/2 \rfloor$ new nodes. This increases the number of nodes in G' and the number of edges by a multiplicative g factor, does not decrease the girth, and ensures the diameter is $O(g)$.

The lower bound graph G consists of G', the node u, the nodes and paths connecting each v'_i to u, and for each v'_i another node v_i connected to it by an edge (v_i, v'_i). Let $V_A = \{v_1, \ldots, v_{n'}\}$ and V_B be the set of all other nodes.

The pairwise spanner we construct approximately preserves distances between pairs of nodes in V_A, i.e. $\mathcal{P} \subseteq V_A \times V_A$. Since the girth of G' is $3\alpha + \beta$, if $e' = \{v'_i, v'_j\}$ is an edge in G' and $(v_i, v_j) \in \mathcal{P}$, then any (α, β)-pairwise spanner must contain the edge e': in G, the distance between v_i and v_j is 3; and if there is a path in H connecting v_i and v_j with length at most $3\alpha + \beta$, then this path connects v'_i and v'_j in G' with $3\alpha + \beta - 2$ edges, which together with e' closes a cycle of length $3\alpha + \beta - 1 < g$ in G'. On the other hand, if $(v_i, v_j) \notin \mathcal{P}$ then the edge e' can be omitted from the spanner.

Proof of Theorem 26. Fix a distributed protocol σ for constructing an (α, β)-pairwise spanner. Let G be the graph described above, and denote the edges of G' by $e_1, \ldots, e_m$. Theorem 25 gives the bound $m \geq 2cn^{1+\frac{4}{3g-10}}$ for some constant $c > 0$.

We use the same reduction from $\mathsf{PART\text{-}COMP}_{m,p}$ to σ: Alice has a set $x \subseteq \{1, \ldots, m\}$ of size p, and Bob has to output a set $y \subseteq \{1, \ldots, m\}$ of size $m/2$ satisfying $x \cap y = \emptyset$. They simulate σ on the graph G and the set of pairs $\mathcal{P} = \left\{(v_i, v_j) : \exists k \in x \; e_k = \{v'_i, v'_j\}\right\}$; Alice simulates the nodes in V_A and Bob simulates the nodes V_B and the edges among them. To simulate communication on edges of the form (v_i, v'_i), Alice and Bob communicate. The bound on p in the theorem statement comes from the fact that $|\mathcal{P}| \leq m/2$.

The spanner constructed is a subgraph H of G with at most $m/2$ edges, satisfying $\delta_H(v_i, v_j) \leq \alpha \cdot \delta_G(v_i, v_j) + \beta$ for all $(v_i, v_j) \in \mathcal{P}$. For each such pair, the edge $\{v'_i, v'_j\}$ must be in H, as explained above. Let $y = \{k : e_k \in E_G \setminus E_H\}$. The spanner size implies $|y| \geq m/2$, while the above discussion implies $x \cap y = \emptyset$. Thus, Bob can output a subset of y of size $m/2$, solving the communication complexity problem.

By the communication complexity lower bound, Alice and Bob must communicate $\Omega(p)$ bits during the simulation. The number of edges they simulate together is at most n, and $O(\log n)$ bits are transferred over each edge at each round. Thus, the protocol must take $\Omega\left(\frac{|\mathcal{P}|}{n \log n}\right)$ rounds to complete. □

5 Discussion

This paper presents various algorithms for computing sparse purely additive spanners in the CONGEST model. Our algorithms exhibit tradeoffs between the running time and the sparsity of the constructed spanners. By choosing different values for the parameter h, one can obtain a spanner with the same stretch in a smaller number of rounds but at the expense of increasing the density. This tradeoff is an important direction for future work.

Our lower bound uses a new communication complexity problem, and leverages the distributed nature of the system by using the fact that each node initially only knows the pairs in $\mathcal{P}$ to which it belongs[3]. That is, the topology of the graph used for the lower bound reduction is known completely to both Alice and Bob, regardless of their inputs to the $\mathsf{PART\text{-}COMP}_{m,p}$ instance, while the uncertainty about the identity of the pairs in $\mathcal{P}$ is what makes the problem hard. While it might be unnatural to assume that other nodes know about these pairs, it is theoretically interesting to ask whether one can design faster distributed constructions given this information.

Finally, we believe that our new lower bound technique can be useful for proving additional lower bounds in the CONGEST model, as it diverges from reducing to the set-disjointness problem.

Acknowledgements We thank Merav Parter for a helpful discussion on the lower bound, and the anonymous referees of DISC 2016 for valuable comments.

References

[AB16a] Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. In *ACM SIGACT Symposium on Theory of Computing, STOC*, 2016.

[AB16b] Amir Abboud and Greg Bodwin. Error amplification for pairwise spanner lower bounds. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 841–854, 2016.

[ACIM99] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.

[ADD+93] Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

[Bas08] Surender Baswana. Streaming algorithm for graph spanners - single pass and constant processing time per edge. *Inf. Process. Lett.*, 106(3):110–114, 2008.

[BCE05] Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. *SIAM J. Discrete Math.*, 19(4):1029–1055, 2005.

[BKMP10] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (alpha, beta)-spanners. *ACM Trans. Algorithms*, 7(1):5, 2010.

[BKS12] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. Fully dynamic randomized algorithms for graph spanners. *ACM Trans. Algorithms*, 8(4):35, 2012.

[3] In fact, our lower bound holds even if all nodes in pairs in $\mathcal{P}$ know all of $\mathcal{P}$.

[BS07] Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007.

[BS08] Surender Baswana and Soumojit Sarkar. Fully dynamic algorithm for graph spanners with poly-logarithmic update time. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1125–1134, 2008.

[BW16] Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 855–872, 2016.

[CE06] Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discrete Math.*, 20(2):463–501, 2006.

[CGK13] Marek Cygan, Fabrizio Grandoni, and Telikepalli Kavitha. On pairwise spanners. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS*, pages 209–220, 2013.

[CGK14] Keren Censor-Hillel, Mohsen Ghaffari, and Fabian Kuhn. Distributed connectivity decomposition. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 156–165, 2014.

[Che13a] Shiri Chechik. Compact routing schemes with improved stretch. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 33–41, 2013.

[Che13b] Shiri Chechik. New additive spanners. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 498–512, 2013.

[CHKM12] Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC*, pages 961–970, 2012.

[CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.

[DG08] Bilel Derbel and Cyril Gavoille. Fast deterministic distributed algorithms for sparse spanners. *Theor. Comput. Sci.*, 399(1-2):83–100, 2008.

[DGP07] Bilel Derbel, Cyril Gavoille, and David Peleg. Deterministic distributed construction of linear stretch spanners in polylogarithmic time. In *Distributed Computing, 21st International Symposium, DISC*, pages 179–192, 2007.

[DGPV08] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. On the locality of distributed sparse spanner construction. In *Proceedings of the 27th Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 273–282, 2008.

[DGPV09] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. Local computation of nearly additive spanners. In *Distributed Computing, 23rd International Symposium, DISC*, pages 176–190, 2009.

[DHZ00] Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. *SIAM J. Comput.*, 29(5):1740–1759, 2000.

[DKO14] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 367–376, 2014.

[DMP+05] Devdatt P. Dubhashi, Alessandro Mei, Alessandro Panconesi, Jaikumar Radhakrishnan, and Aravind Srinivasan. Fast distributed algorithms for (weakly) connected dominating sets and linear-size skeletons. *J. Comput. Syst. Sci.*, 71(4):467–479, 2005.

[DSHK+12] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.

[Elk05] Michael Elkin. Computing almost shortest paths. *ACM Trans. Algorithms*, 1(2):283–323, 2005.

[Elk07a] Michael Elkin. A near-optimal distributed fully dynamic algorithm for maintaining sparse spanners. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC*, pages 185–194, 2007.

[Elk07b] Michael Elkin. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. In *Automata, Languages and Programming, 34th International Colloquium, ICALP*, pages 716–727, 2007.

[EP04] Michael Elkin and David Peleg. (1+epsilon, beta)-spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.

[Erd64] P. Erdős. Extremal problems in graph theory. In *Theory of graphs and its applications: proceedings of the symposium held in Smolenice in June 1963*, pages 29–36. Pub. House of the Czechoslovak Academy of Sciences, 1964.

[EZ06] Michael Elkin and Jian Zhang. Efficient algorithms for constructing (1+epsilon, beta)-spanners in the distributed and streaming models. *Distributed Computing*, 18(5):375–385, 2006.

[FHW12] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1150–1162, 2012.

[GK13] Mohsen Ghaffari and Fabian Kuhn. Distributed minimum cut approximation. In *Distributed Computing - 27th International Symposium, DISC*, pages 1–15, 2013.

[HP14] Stephan Holzer and Nathan Pinsker. Approximation of distances and shortest paths in the broadcast congest clique. *CoRR*, abs/1412.3445, 2014.

[HW12] Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 355–364, 2012.

[Kav15] Telikepalli Kavitha. New pairwise spanners. In *32nd International Symposium on Theoretical Aspects of Computer Science, STACS*, pages 513–526, 2015.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.

[Knu14] Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. In *14th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT*, pages 277–281, 2014.

[KV13] Telikepalli Kavitha and Nithin M. Varma. Small stretch pairwise spanners. In *40th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 601–612, 2013.

[LP13] Christoph Lenzen and David Peleg. Efficient distributed source detection with limited bandwidth. In *ACM Symposium on Principles of Distributed Computing, PODC*, pages 375–382, 2013.

[Mat02] Jiri Matousek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

[Par14] Merav Parter. Bypassing erdős' girth conjecture: Hybrid stretch and sourcewise spanners. In *41st International Colloquium on Automata, Languages, and Programming, ICALP*, pages 608–619, 2014.

[Pel00] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, 2000.

[Pet09] Seth Pettie. Low distortion spanners. *ACM Trans. Algorithms*, 6(1), 2009.

[Pet10] Seth Pettie. Distributed algorithms for ultrasparse spanners and linear size skeletons. *Distributed Computing*, 22(3):147–166, 2010.

[PR99] David Peleg and Vitaly Rubinovich. A near-tight lower bound on the time complexity of distributed MST construction. In *40th Annual Symposium on Foundations of Computer Science, FOCS*, pages 253–261, 1999.

[PS89] David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.

[PU89a] David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.

[PU89b] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.

[RTZ05] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *32nd International Colloquium on Automata, Languages and Programming, ICALP*, pages 261–272, 2005.

[RZ11] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011.

[TZ01] Mikkel Thorup and Uri Zwick. Compact routing schemes. In *SPAA*, pages 1–10, 2001.

[TZ05] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52(1), 2005.

[TZ06] Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 802–809, 2006.

[Woo10] David P. Woodruff. Additive spanners in nearly quadratic time. In *37th International Colloquium on Automata, Languages and Programming, ICALP*, pages 463–474, 2010.